To appear 1995 IEEE Int'l Conf. on Systems, Man and Cybernetics, October 1995 #15

Decision Problems Related to Structural Induction for Rings of Petri Nets with Fairness*

Jianan Li † Ichiro Suzuki † Masafumi Yamashita ‡
†Department of Electrical Engineering and Computer Science
University of Wisconsin–Milwaukee, Milwaukee, WI 53201, U.S.A.
‡Department of Electrical Engineering, Faculty of Engineering
Hiroshima University, Higashi-Hiroshima 724, Japan

Abstract Structural induction is a technique for proving that a system consisting of many identical components works correctly regardless of the actual number of components it has. Previously the authors have obtained conditions under which structural induction goes through for rings that are modeled as a Petri net satisfying a fairness requirement. The conditions guarantee that for some k, all rings of size k or greater exhibit "similar" behavior. The key concept is the similarity between rings, where rings R^k and R^ℓ of sizes k and ℓ, respectively, are said to be similar if, intuitively, (1) none of the components in either ring can tell whether it is in R^k or R^ℓ, and (2) none of the components (except possibly one) can tell its position within the ring to which it belongs. A ring satisfying this second property is said to be *uniform*. In this paper we prove the undecidability of various basic questions regarding similarity and uniformity. Some of the questions shown to be undecidable are: (1) Is there k such that R^k and R^{k+1} are similar? (2) Is there k such that all rings of size k or greater are mutually similar? (3) Is there k such that R^k is uniform? (4) Is there k such that $R^k, R^{k+1}, R^{k+2}, \dots$ are all uniform?

1 Introduction

Given a system consisting of many identical finite state components that are connected in some regular topology, how can we determine whether it works correctly regardless of the actual number of components it has? Conventional theorem provers based on state-space search cannot be used directly to answer this question, since infinitely many instances of the system are involved. In fact, this problem is known to be unsolvable in general, even if the topology of the system is a unidirectional ring [1] [10]. The works reported in [2] [3] [5] [6] [7] [13] are some of the efforts to find a sufficient condition for such a system to be correct regardless of the number of components.

*This work was supported in part by the National Science Foundation under grants CCR-9004346 and IRI-9307506, the Office of Naval Research under grant N00014-94-1-0284, and an endowed chair (KIFUKOZA) supported by Hitachi Ltd. at Faculty of Engineering Science, Osaka University.

Inspired by [5] [13], the authors have considered the analysis problem stated above for *rings* of identical components given as a Petri net satisfying a fairness requirement, and obtained structural induction theorems that can be used to formally infer the correctness of a ring of any large size from the correctness of a ring having only a few components [6] [7]. Petri nets (see, for example, [9]) are widely used for modeling and analysis of concurrent processing systems. Fairness is an important property of concurrent systems often studied in the context of temporal logic [8]. The combination of Petri nets and temporal logic has been found to be extremely useful for formal analysis of such systems [4] [11] [12].

One of the key concepts in the authors' induction theorems is the "similarity" between two rings. For $k \geq 2$, let R^k be the ring consisting of k identical copies $C_0, C_1, \dots, C_{k-1}$ of a component C. We assume that all components except possibly C_0 have the same initial state. Intuitively, we say that rings R^k and R^ℓ are *similar* if (1) none of the components in either ring can tell whether it is in R^k or R^ℓ, and (2) none of the components (except possibly C_0) can tell its position within the ring to which it belongs. (A ring satisfying this second property is said to be *uniform*.) The induction theorems reported in [6][7] provide sufficient conditions for all rings of size k or greater to be mutually similar for some k. Note that if R^k is correct in some sense and all rings of size k or greater are mutually similar, then we can conclude that R^ℓ is correct for all $\ell \geq k$. The theorems have been applied to the formal analysis of token-passing mutual exclusion, a simple producer-consumers system, and demand-driven token-circulation.

The goal of this paper is to prove the undecidability of the following basic questions regarding similarity and uniformity (as well as "weak similarity" defined later):

1. Is there k such that R^k and R^{k+1} are similar?
2. Is there k such that all rings of size k or greater are mutually similar?
3. Is there $k \geq 3$ such that R^k is uniform?
4. Are $R^3, R^4, \dots$ all uniform?

JGK

5. Is there $k \geq 3$ such that $R^k, R^{k+1}, R^{k+2}, \ldots$ are all uniform?

These negative results might seem somewhat expected in view of the results in [1] [10] and the undecidability results reported in [13] regarding the existence of "network invariants" that are needed for carrying out induction. However, similarity and uniformity of rings are stronger requirements than the existence of network invariants, and hence the proofs of the undecidability results reported here require certain unique arguments.

As is noted in [13] for the invariant method, despite these negative results we still expect that for many interesting, practical ring systems, the induction theorems of [6] [7] can be an effective tool for formal analysis.

The rest of the paper is organized as follows. Section 2 reviews the basic terminology of Petri nets. Section 3 introduces the basic concepts regarding components and rings. In Section 4 we introduce similarity and related concepts. The undecidability results are presented in Section 5. The concluding remarks are found in Section 6.

2 Petri Nets

We review the standard terminology of Petri nets [9].

A *Petri net* is a directed graph with two types of nodes, called *transitions* and *places*, and weighted arcs from a node of one type to a node of the other type. Formally, it is given as a triple $N = (P, T, F)$, where P is a finite set of places, T is a finite set of transitions, and $F : (P \times T) \cup (T \times P) \rightarrow \{0, 1, 2, \ldots\}$ is a weight function. A place $p \in P$ is called an *input place* (or *output place*) of a transition $t \in T$ if $F(p,t) \geq 1$ (or $F(t,p) \geq 1$). Any function $M : P \rightarrow \{0, 1, 2, \ldots\}$ is called a *marking*. A place p is said to have $M(p)$ *tokens* at a marking M. A transition $t \in T$ is said to be *firable* at M iff $M(p) \geq F(p,t)$ for every $p \in P$. If t is firable at M, then it may *fire* and yield another marking M' such that $M'(p) = M(p) - F(p,t) + F(t,p)$ for every $p \in P$. We denote this by $M \rightarrow_t M'$. This relation is extended by

1. $M \rightarrow_\lambda M$ and
2. $M \rightarrow_{\alpha t} M'$ iff there exists M'' such that $M \rightarrow_\alpha M''$ and $M'' \rightarrow_t M'$

for all M, M', $\alpha \in T^*$ and $t \in T$.[1] If $M \rightarrow_\alpha M'$ then M' is said to be *reachable* from M by a *finite firing sequence* α. $L(N, M)$ denotes the set of all finite firing sequences from M. An infinite sequence $\alpha \in T^\omega$ is an *infinite firing sequence* from M if $\beta \in L(N, M)$ for every finite prefix β of α. We denote by $L^\omega(N, M)$ the set of infinite firing sequences from M. Let $L^\infty(N, M) = L(N, M) \cup L^\omega(N, M)$ denote the set of *all* (both finite and infinite) firing sequences from M. Usually an *initial marking* is associated with a Petri net.

We draw a Petri net using a circle and a square to represent places and transitions, respectively. An arc with weight $F(p,t)$ (or $F(t,p)$) is drawn from p to t (or from t to p) if $F(p,t) \geq 1$ (or $F(t,p) \geq 1$). The weight is omitted if it is 1. A marking M is represented by drawing $M(p)$ dots in (the circle representing) p.

When describing a system using a Petri net $N = (P, T, F)$, we may designate a subset $T' \subseteq T$ of transitions such that every transition $t \in T'$ must be fired *fairly*, i.e., if t becomes firable infinitely often, then it must fire infinitely often. Furthermore, to examine whether net N with initial marking M has certain liveness or eventuality properties (e.g., "if t_1 fires then eventually t_2 fires"), we may only be interested in those firing sequences α that are either infinite, or finite and *terminating* in the sense that there is no transition t such that $\alpha t \in L(N, M)$. This observation leads us to use the set $\mathcal{L}(N, M, T')$ defined below to examine the behavior of N: $\mathcal{L}(N, M, T')$ is the set of all firing sequences $\alpha \in L^\infty(N, M)$ such that either

1. α is infinite and satisfies the fairness requirement on the transitions in T', or
2. α is finite and terminating.

3 Components and Rings

To save space, we introduce the necessary concepts informally, using examples. The reader is referred to [6] [7] for a formal discussion.

A *component* is a Petri net $C = (P, T, F)$ in which the set T of transitions can be partitioned into three groups, the *left interface transitions*, the *internal transitions* and the *right interface transitions*, where the number of left interface transitions must equal the number of right interface transitions. See Figure 1(a) for an example of a component having two left interface transitions u_1 and u_2, three internal transitions v_1, v_2 and v_3, and two right interface transitions w_1 and w_2. We connect two or more copies of C to form a ring, as shown in Figure 1(b), by merging the respective interface transitions of adjacent copies of C. For any $k \geq 2$, the ring consisting of k copies of C is denoted R^k, and the copies of C in R^k are referred to as C_0, C_1, ..., C_{k-1}, where C_{i+1} is the right neighbor of C_i. (Subscripts are taken modulo k when we discuss R^k). Formally, the transitions and places of C must be renamed in each C_i (e.g., v_2 of C_1 might be renamed "$v_{1,2}$" in the example given above) so that they all have distinct names in R^k. However, for convenience, in Figure 1(b)

[1] T^* and T^ω denote, respectively, the set of finite sequences and the set of infinite sequences of the elements of T. λ is the empty sequence.

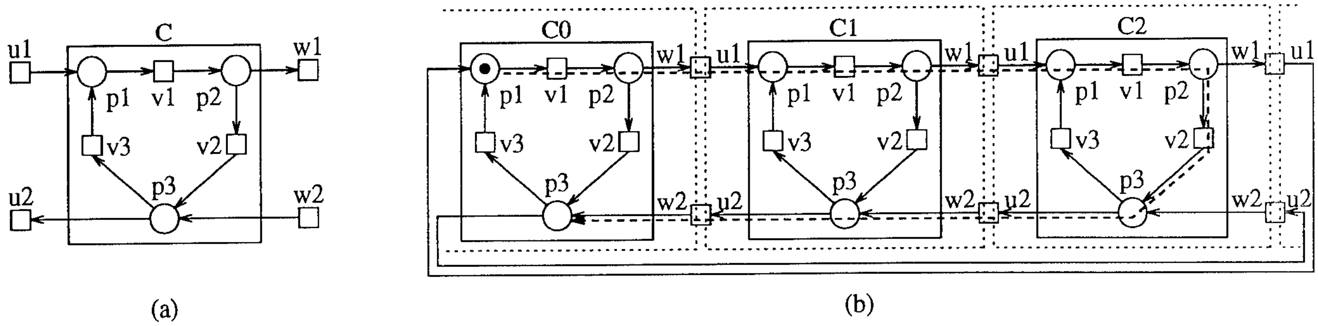

Figure 1: (a) Component C. (b) Ring R^3 consisting of three copies of C.

we use their original names in each copy of C, and thus assign two names to every interface transition.

We assume that for all $k \geq 2$, the initial marking M^k of R^k is such that for some fixed markings M and M' of C, C_0 has M and all other C_i's, $i \geq 1$, have M'. (It is often necessary to break symmetry by giving C_0 an initial marking different from that of the other components.) Furthermore, we assume that for all $k \geq 2$, fairness (when required) should be imposed on an identical set of transitions in all components of R^k.

In the following, sets $L^\infty(R^k, M^k)$ and $\mathcal{L}(R^k, M^k, T^k)$ are simply written as $L^\infty(k)$ and $\mathcal{L}(k)$, respectively, where T^k the set of transitions of R^k (consisting of identical selections from all components) on which fairness is required.

Let $\alpha \in L^\infty(k)$ be a firing sequence of R^k. Let C_i be a component in R^k. We define the *local history* of C_i in α, denoted $\langle\langle C_i\rangle\rangle_\alpha$, to be the sequence obtained from α by deleting all transitions except those that belong to C_i. Here, a transition t *belongs to* C_i iff t is either a left interface transition, a right interface transition, or an internal transition of C_i. (So, an interface transition between C_i and C_{i+1} belongs to both C_i and C_{i+1}.) Sequence $\langle\langle C_i\rangle\rangle_\alpha$ is thus the portion of α that occurs in C_i.

As we mentioned above, formally in R^k the transitions (and places) are renamed so that all transitions have distinct names. Therefore, to be able to compare local histories of different components (that may even belong to different rings R^k and R^ℓ), when describing $\langle\langle C_i\rangle\rangle_\alpha$ we use the *original* transition names given in C, rather than the names in R^k after the renaming (e.g., use "v_2" instead of "$v_{1,2}$" when describing $\langle\langle C_1\rangle\rangle_\alpha$).

Example 1 Consider ring R^3 shown in Figure 1(b). Suppose that in the initial marking, only C_0 has a token, in place p_1, as shown in the figure. Consider the firing sequence α that moves the token through C_1 and C_2 and back to C_0, as indicated by the dashed arrow in Figure 1(b). Then the local history of C_1 in α is $\langle\langle C_1\rangle\rangle_\alpha = u_1v_1w_1w_2u_2$. Similarly, the reader can verify $\langle\langle C_0\rangle\rangle_\alpha = v_1w_1w_2$ and $\langle\langle C_2\rangle\rangle_\alpha = u_1v_1v_2u_2$. □

Let M be a marking of R^k. The firability of an interface transition t of C_i at M is determined by the token counts of the input places of t, and such places may belong to C_{i-1}, C_i or C_{i+1}. We define the *firability vector* of C_i at M to be a column vector of token counts of *all* input places (in C_{i-1}, C_i and C_{i+1}) of *all* interface transitions of C_i. (The order in which the token counts of these places appear must be the same for all components.)

Let α be a firing sequence of R^k. Note that during the execution of α, a firing of a transition in C_{i-1}, C_i and C_{i+1} may change the firability vector of C_i. To describe the changes in the firability vector of C_i in α, we now define the *extended local history* of C_i in α, denoted $\langle C_i\rangle_\alpha$, as follows: $\langle C_i\rangle_\alpha$ is obtained from $\langle\langle C_i\rangle\rangle_\alpha$ by inserting the firability vector of C_i at the corresponding positions each time it changes. Again, instead of a formal definition we give an example.

Example 2 Consider firing sequence α discussed in the previous example. The reader can verify that the extended local history $\langle C_1\rangle_\alpha$ of C_1 in α is

$$\begin{bmatrix}0\\0\\0\\0\end{bmatrix} \begin{bmatrix}1\\0\\0\\0\end{bmatrix} u_1 \begin{bmatrix}0\\0\\0\\0\end{bmatrix} v_1 \begin{bmatrix}0\\0\\1\\0\end{bmatrix} w_1 \begin{bmatrix}0\\0\\0\\0\end{bmatrix} \begin{bmatrix}0\\0\\0\\1\end{bmatrix} w_2 \begin{bmatrix}0\\1\\0\\0\end{bmatrix} u_2 \begin{bmatrix}0\\0\\0\\0\end{bmatrix}$$

where the vectors show, in the following order, the token counts of p_2 of C_0 (the input place of u_1 of C_1), p_3 of C_1 (the input place of u_2 of C_1), p_2 of C_1 (the input place of w_1 of C_1), and p_3 of C_2 (the input place of w_2 of C_1). □

Intuitively, $\langle\langle C_i\rangle\rangle_\alpha$ describes not only which transitions of C_i fire, but also how the "environment" of C_i changes in α in terms of the firability of its interface transitions. Of course, if $\langle\langle C_i\rangle\rangle_\alpha = \langle\langle C_i\rangle\rangle_\beta$ for firing sequences α and β, then C_i cannot distinguish α and β. On the other hand, the information on the behavior of the "environment" of C_i is not included in local histories $\langle C_i\rangle_\alpha$ and $\langle C_i\rangle_\beta$.

4 Similarity and Uniformity

We now review the main concepts introduced in [6] [7] for stating the induction theorems. Recall that $\mathcal{L}(k)$ and $\mathcal{L}(\ell)$ are the set of firing sequences of ring R^k and R^ℓ, respectively, that we use to determine their correctness.

Definition 1 [6] Two rings R^k and R^ℓ are *similar*, denoted $R^k \sim R^\ell$, if

1. $\{\langle C_0\rangle_\alpha | \alpha \in \mathcal{L}(k)\} = \{\langle C_0\rangle_\alpha | \alpha \in \mathcal{L}(\ell)\}$ and
2. $\{\langle C_i\rangle_\alpha | \alpha \in \mathcal{L}(k)\} = \{\langle C_j\rangle_\alpha | \alpha \in \mathcal{L}(\ell)\}$ for any $1 \le i \le k-1$ and $1 \le j \le \ell - 1$.

Definition 2 [7] Two rings R^k and R^ℓ are *weakly similar*, denoted $R^k \overset{\bullet}{\sim} R^\ell$, if

1. $\{\langle\langle C_0\rangle\rangle_\alpha | \alpha \in \mathcal{L}(k)\} = \{\langle\langle C_0\rangle\rangle_\alpha | \alpha \in \mathcal{L}(\ell)\}$ and
2. $\{\langle\langle C_i\rangle\rangle_\alpha | \alpha \in \mathcal{L}(k)\} = \{\langle\langle C_j\rangle\rangle_\alpha | \alpha \in \mathcal{L}(\ell)\}$ for any $1 \le i \le k-1$ and $1 \le j \le \ell - 1$.

Intuitively, if $R^k \sim R^\ell$, then none of the copies of C knows which of R^k and R^ℓ it is in, and none of the copies of C other than C_0 knows which copy of C it is. (C_0 might behave differently from others, since its initial marking may not be the same as that of others.) Weak similarity is identical to similarity, except that the firability vectors of components are not considered. By definition, $R^k \sim R^\ell$ implies $R^k \overset{\bullet}{\sim} R^\ell$.

As is explained in [6][7], one way to prove the correctness of R^n for all values of $n \ge k$ for some k is to show that

1. $R^k \overset{\bullet}{\sim} R^{k+1} \overset{\bullet}{\sim} R^{k+2} \overset{\bullet}{\sim} \cdots$,
2. $S \subseteq \{\langle\langle C_0\rangle\rangle_\alpha | \alpha \in \mathcal{L}(k)\} \subseteq S'$, where S and S' are sets of firing sequences of C describing the correctness requirements for C_0, and
3. $S'' \subseteq \{\langle\langle C_j\rangle\rangle_\alpha | \alpha \in \mathcal{L}(k)\} \subseteq S'''$ for *any one* j, $1 \le j \le k-1$, where S'' and S''' are sets of firing sequences of C describing the correctness requirements for all C_i, $i \ge 1$.

Then C_0's in all $R^k, R^{k+1}, \ldots$ have identical properties and C_0 of R^k is correct. Thus C_0 is correct in all $R^k, R^{k+1}, \ldots$. Similarly, all C_i's, $i \ge 1$, in all $R^k, R^{k+1}, \ldots$ have identical properties and that particular C_j of R^k (mentioned above) is correct. So all these C_i's in all $R^k, R^{k+1}, \ldots$ are also correct. The main result in [6] is a sufficient condition for $R^2 \sim R^3 \sim R^4 \sim \cdots$ (hence $R^2 \overset{\bullet}{\sim} R^3 \overset{\bullet}{\sim} R^4 \overset{\bullet}{\sim} \cdots$) to hold, and [7] presents a sufficient condition for $R^k \overset{\bullet}{\sim} R^{k+1} \overset{\bullet}{\sim} R^{k+2} \overset{\bullet}{\sim} \cdots$ to hold for the given $k \ge 5$. (We do not review these sufficient conditions since doing so requires additional definitions. The interested reader is referred to [6][7].) The concept of "uniformity" introduced next is used to state the condition in [7].

Definition 3 [7] R^k, $k \ge 3$, is *uniform* if $\{\langle\langle C_i\rangle\rangle_\alpha | \alpha \in \mathcal{L}(k)\} = \{\langle\langle C_j\rangle\rangle_\alpha | \alpha \in \mathcal{L}(k)\}$ for any $1 \le i, j \le k-1$, $i \ne j$.

Note that if R^k is uniform, then none of the components $C_1, C_2, \ldots, C_{k-1}$ can determine its position in the ring.

Example 3 We can show that the "environment" looks identical to all C_i, $i \ge 1$, in any R^k, $k \ge 3$, constructed from C of Figure 1(a). On the other hand, R^2 and R^3 are slightly different. In R^2, w_2 of C_1 becomes firable when C_1 fires u_2, since its left and right neighbors are both C_0. This cannot happen in R^3, since the left and right neighbors of C_1 are different. If we ignore the firability vectors, of course, R^2 and R^3 become indistinguishable to C_1. Using an argument along this line, we can prove $R^2 \not\sim R^3 \sim R^4 \sim \cdots$, $R^2 \overset{\bullet}{\sim} R^3 \overset{\bullet}{\sim} R^4 \overset{\bullet}{\sim} \cdots$, and that R^k is uniform for all $k \ge 3$. $\square$

In [6][7], the argument outlined above has been used to prove the correctness of rings for token-passing mutual exclusion, a simple producer-consumers system, and demand-driven token-circulation.

5 Undecidability Results

As we discussed in Section 4, similarity, weak similarity and uniformity can be a basis for proving that a ring system is correct regardless of its size. In this section we prove that the basic questions regarding these concepts posed in Section 1 are undecidable in general.

Theorem 1 *Given C, M and M', the following problems are undecidable in general, even if fairness is not required on any transition.*

1. *Is there k such that $R^k \sim R^{k+1}$?*
2. *Is there k such that $R^k \overset{\bullet}{\sim} R^{k+1}$?*
3. *Is there k such that $R^k \sim R^{k+1} \sim R^{k+2} \sim \cdots$?*
4. *Is there k such that $R^k \overset{\bullet}{\sim} R^{k+1} \overset{\bullet}{\sim} R^{k+2} \overset{\bullet}{\sim} \cdots$?*

Proof We first prove the second and fourth claims that do not involve the changes in the firability vectors of the components. The basic idea is to simulate the given Turing machine $\mathcal{A}$ (that does not halt in two steps) on the semi-infinite blank tape for (up to) n steps using a uniform ring R^n of size n, in such a way that (1) if $\mathcal{A}$ halts

in n steps, then $R^n \stackrel{\bullet}{\sim} R^{n+1} \stackrel{\bullet}{\sim} R^{n+2} \stackrel{\bullet}{\sim} \cdots$, and (2) if $\mathcal{A}$ never halts, then $R^k \stackrel{\bullet}{\not\sim} R^\ell$ for all $k \neq \ell$. In the sense that we use R^n to simulate a Turing machine or a two-counter automaton for n steps, the proof is similar to those found in [10] and [14]. But here the proof is technically more involved, since weak similarity is a fairly strong condition.

Given a Turing machine $\mathcal{A}$, we construct a component C and markings M and M' of C such that R^n simulates the computation of $\mathcal{A}$ on the semi-infinite blank tape for (up to) n steps. We will first describe the components as a finite state machine, and then later explain how we can represent them by a Petri net. Number the tape cells $0, 1, \ldots$ starting with the leftmost one. For each $0 \leq i \leq n-1$, C_i maintains cell i. Initially, the cells are all blank, and C_0 has the tape head (i.e., the tape head is reading cell 0). The component that has the tape head also remembers in its finite control the current state of $\mathcal{A}$. To simulate a single step of $\mathcal{A}$, the component, say C_i, that currently has the tape head (1) rewrites the symbol in cell i, (2) determines the next state q of $\mathcal{A}$, and then (3) sends q to either C_{i-1} or C_{i+1} depending on whether the tape head is moved left or right. The neighbor of C_i that receives q knows that it now has the tape head and simulates the next step of $\mathcal{A}$.

To ensure that the simulation stops in n steps and we do not run out of space (R^n has exactly n tape cells), we use a flag COUNT (initially false) in each component and two special symbols called NEWSTEP and EXECUTE in the following manner. Intuitively, we set one COUNT in the ring to true before each step of simulation. Specifically, suppose C_i wishes to simulate a step of $\mathcal{A}$. If its own COUNT is still false, then it sets COUNT to true and sends EXECUTE to the right and waits for EXECUTE to arrive from the left. If its own COUNT is true, then it sends NEWSTEP to the right and waits for EXECUTE to arrive from the left. In either case, C_i simulates a step of $\mathcal{A}$ as described above when EXECUTE arrives. A component that does not have the tape head always passes EXECUTE to the right. A component that receives NEWSTEP for the first time (in this case COUNT is still false) changes its COUNT to true and sends EXECUTE (instead of NEWSTEP) to the right. A component that receives NEWSTEP when its COUNT is true passes NEWSTEP to the right.

The simulation ends when either (1) component C_i that has the tape head simulates one step of $\mathcal{A}$ and finds that the next state is a halting state, or (2) C_0 receives either NEWSTEP or a state of $\mathcal{A}$ from the left. In the first case, $\mathcal{A}$ has halted within n steps, and C_i circulates a special symbol HALT once around the ring and halts. In the second case, $\mathcal{A}$ did not halt within n steps, and C_0 circulates a special symbol SUSPEND once around the ring and halts. In either case, the entire ring comes to a halt after the circulation of HALT or SUSPEND.

Since each component has only a finite number of states, constructing C as a Petri net is straightforward. Transfers of various symbols (NEWSTEP, EXECUTE, HALT, SUSPEND and a state of $\mathcal{A}$) can be represented by different interface transitions. The fact that initially C_0 has the tape head can be reflected in its initial marking M different from M' of other components. Fairness is not imposed on any transition.

Now, we modify C slightly to make R^n uniform. That is, before the simulation starts, we give any component in R^n a chance to "act" as C_0 in the simulation. We achieve this by letting (the real) C_0 nondeterministically do one of the following: (1) start the simulation, (2) send a token to the left through a special interface transition, and (3) send a token to the right through another special interface transition. A component that receives a token through one of the special interface transitions can, nondeterministically, either start the simulation acting as C_0 (as described above) or simply pass the token to the neighbor on the opposite side. This modification guarantees that R^n is uniform. A side effect is that the simulation of $\mathcal{A}$ may never start, but, this only adds two (identical) sequences (one for each direction in which the token is passed forever) to the sets of local histories of all C_i's, $i \geq 1$, and two sequences to the set of local histories of C_0. So whether or not two components have identical sets of local histories is not affected.

Suppose that $\mathcal{A}$ never halts. Then for any n, (the component acting as) C_0 of R^n receives EXECUTE exactly n times. Thus for any distinct k and ℓ, the two conditions of Definition 2 do not hold, and hence $R^k \stackrel{\bullet}{\not\sim} R^\ell$.

Suppose that $\mathcal{A}$ halts within $n \geq 3$ steps. Then for any $k \geq n$, in R^k, (the component acting as) C_0 and its $n-1$ right neighbors (call them $C_1, \ldots, C_{n-1}$ for convenience) perform the actual simulation (the simulation itself does not depend on the value of k) and each of $C_n, \ldots, C_{k-1}$ simply pass EXECUTE n times. (In addition, all components pass HALT once.) This, together with the fact that all rings are uniform, implies that for any $k, \ell \geq n$, the two conditions of Definition 2 are satisfied, and hence $R^k \stackrel{\bullet}{\sim} R^\ell$. This completes the proof of the second and fourth claims.

As for the first and third claims of the lemma, it is easy to see that component C described above can be constructed in such a way that each interface transition t has exactly one input place, say p, and once p receives a token it loses the token only when t fires. Then given any firing sequence α of R^k, the changes in the firability vectors during the execution of α can be determined completely. So the conditions of Definition 1 are satisfied iff those of Definition 2 are satisfied. Therefore the claims we made on the relation between the behavior of $\mathcal{A}$ and weak similarity among rings hold also for similarity among rings. This completes the proof of the first and third claims. □

Now we discuss uniformity.

Theorem 2 *Given C, M and M', the following problems are undecidable in general, even if fairness is not required on any transition.*

1. *Is there $k \geq 3$ such that R^k is uniform?*

2. *Are $R^3, R^4, \ldots$ all uniform?*

3. *Is there $k \geq 3$ such that $R^k, R^{k+1}, R^{k+2}, \ldots$ are all uniform?*

Proof We modify the construction given in the proof of Theorem 1. Note that the real C_0 can transfer its role as the initiator of the simulation to other components, but even after the transfer, it can still "remember" that it is the real C_0. So, when the real C_0 receives SUSPEND, (in addition to passing it as before) we let it fire a special right interface transition. Then the local history of the real C_1 (the right neighbor of C_0) becomes different from those of any other component, since that interface transition never fires at other components. So if $\mathcal{A}$ never halts, then the special transition is always fired, and thus R^n is not uniform for any n. On the other hand, if $\mathcal{A}$ halts within $n \geq 3$ steps, then R^k is uniform for all $k \geq n$. This completes the proof of the first claim. To prove the second and third claims, all we need to do is to change the behavior of the real C_0 so that it fires that special right interface transition when it receives HALT, instead of SUSPEND. Then if $\mathcal{A}$ halts in $n \geq 3$ steps, then R^k is not uniform for any $k \geq n$. If on the other hand $\mathcal{A}$ never halts, then that special transition is never fired, and thus R^n is uniform for all n. This completes the proof of the second and third claims. □

6 Concluding Remarks

Since any finite alphabet can be encoded as binary strings, the negative results presented above remains true even if a component has only four interface transitions on each side, two for sending symbols and two for receiving symbols. It would be interesting to investigate the decidability of analogous problems for the case when the components have exactly one interface transition on each side, and for some restricted classes of Petri nets.

References

[1] K. R. Apt and D. C. Kozen, "Limits for automatic verification of finite-state concurrent systems," *Information Processing Letters 15*, May 1986, pp. 307-309.

[2] E. M. Clarke, O. Grümberg and M. C. Browne, "Reasoning about networks with many identical finite-state processes," *Proceedings of the 5th Annual ACM Symposium on Principles of Distributed Computing*, Calgary, Alberta, Canada, August 1986, pp. 240–248.

[3] S. German and A. P. Sistla, "Reasoning about systems with many processes," *Journal of the ACM 39*, No. 3, 1992, pp. 675–735.

[4] R. R. Howell, L. E. Rosier and H. C. Yen, "A taxonomy of fairness and temporal logic problem for Petri nets," in *Proceedings of the 13th Symposium on Mathematical Foundations of Computer Science 1988*, Lecture Notes in Computer Science 324, Springer-Verlag, New York, 1988.

[5] R. P. Kurshan and K. McMillan, "A structural induction theorem for processes," *Proc. 8th ACM Symp. Principles of Distributed Computing*, Edmonton, Alberta, August 1989, pp. 239–247.

[6] J. Li, I. Suzuki and M. Yamashita, "Fair Petri nets and structural induction for rings of processes," *Theoretical Computer Science*, Vol. 135, 1994, pp. 377–404.

[7] J. Li, I. Suzuki and M. Yamashita, "A new structural induction theorem for rings of temporal Petri nets," *IEEE Transactions on Software Engineering 20*, No. 2, February 1994, pp. 115–126.

[8] Z. Manna and A. Pnueli, *The Temporal Logic of Reactive and Concurrent Systems*, Springer-Verlag, New York, 1992.

[9] T. Murata, "Petri nets: Properties, analysis and applications," *Proceedings of the IEEE 77*, No. 4, 1989, pp. 541–580.

[10] I. Suzuki, "Proving properties of a ring of finite state machines," *Information Processing Letters 28*, 1988, pp. 213–214.

[11] I. Suzuki, "Formal analysis of the alternating bit protocol by temporal Petri nets," *IEEE Transactions on Software Engineering 16*, No. 11, November 1990, pp. 1273–1281.

[12] I. Suzuki and H. Lu, "Temporal Petri nets and their application to modeling and analysis of a handshake daisy chain arbiter," *IEEE Transactions on Computers 38*, No. 5, 1989, pp. 696–704.

[13] P. Wolper and V. Lovinfosse, "Verifying properties of large sets of processes with network invariants," in *Automatic Verification Methods for Finite State Systems*, J. Sifakis (ed.), Lecture Notes in Computer Science 407, Springer Verlag, 1990, pp. 68–80.

[14] P. Wolper, private communication.